1. Load the C14 calibration file with

   ```
   IntCal <- read.csv("http://people.reed.edu/~jones/Courses/IntCal.dat")
   ```

   The following code snippet interpolates the calibration curve, giving yearly estimates:

   ```
   attach(IntCal)
   N <- 50001
   cc <- spline(YearBP,C14age,n=N)
   cs <- spline(YearBP,error,n=N)
   CC <- data.frame(YearBP=cc$x,C14age=cc$y,error=cs$y)
   detach()

   names(CC)
   attach(CC) # allows access to CC variables by name
   ```

   You will use the interpolated data frame CC to get posterior densities for a pair of dates.

2. Suppose that we have radiocarbon dates for two objects: 2450 and 2520. Due to their stratigraphically known order, we know that the carbon dates are in the correct order. However, we don't know that their actual ages differ by the same number of years as their carbon ages: all we know is that the second one is older than the first.

   Write an **R** script to generate a sample from the posterior distribution using the Metropolis-Hastings algorithm. You may use a Uniform prior on the set

   $$(Y1, Y2) : Y1 < Y2, \quad Y2 \in 1 : 5000$$

3. Plot the running mean for one of the dates, and use it to select a burn-in set to discard. Using the *acf()* function, select a sampling interval that will give a roughly uncorrelated subsample.

4. Using the *density()* function, estimate and plot the posterior distribution for each age. You may need to tune the bandwidth parameter (bw).

5. Estimate the mean and standard error for each, and give approximate 95% HPD regions for both.